

Nearly Private Information Retrieval^{*}

Amit Chakrabarti and Anna Shubina

Department of Computer Science, Dartmouth College
Hanover, NH 03755, USA
{ac,ashubina}@cs.dartmouth.edu

Abstract. A private information retrieval scheme is a protocol whereby a client obtains a record from a database without the database operators learning anything about which record the client requested. This concept is well studied in the theoretical computer science literature. Here, we study a generalization of this idea where we allow a small amount of information about the client's intent to be leaked.

Despite having relaxed the privacy requirement, we are able to prove three fairly strong lower bounds on such schemes, for various parameter settings. These bounds extend previously known lower bounds in the traditional setting of perfect privacy and, in one case, improve upon the previous best result that handled imperfect privacy.

1 Introduction

Private information retrieval (PIR) schemes have been a substantial focus of theoretical research in computer science, beginning with the highly influential work of Chor, Goldreich, Kushilevitz, and Sudan [4]. In that paper, as in most subsequent work, a *PIR scheme* means a communication protocol that specifies an interaction between a *client* or *user* and one or more database *servers*. The user wishes to obtain a record from the database without the servers learning anything about which record the user seeks.

A clean and concrete version of this problem, as proposed by Chor et al., is as follows: the database y is a string of n bits. The client has an index $j \in \{1, 2, \dots, n\}$ and wishes to obtain the j th bit of y , without the servers obtaining any information about j . As shown by Chor et al., this strong privacy requirement means that if there is only one server that holds the database, the trivial protocol in which the client simply downloads the entire database is optimal in terms of the number of bits communicated. However, as shown in the same paper, if one allows the database to be replicated and copies held by two or more servers that do not talk to each other, the problem can be solved using sublinear communication.

Almost all past work on PIR schemes has required that the servers learn *zero* information about the client's index j . Here, we ask the question: what happens if we allow the protocol to leak a small amount of information about j ? To the best of our knowledge, the only other work to have considered this question is that

^{*} This work was supported in part by an NSF CAREER Award CCF-0448277 and startup funds from Dartmouth College.

of Goldreich, Karloff, Schulman, and Trevisan [5]. It is *a priori* conceivable that relaxing the privacy requirement might decrease the communication required in PIR protocols. However, in this work, we prove three lower bounds that show that previously known lower bounds for traditional (perfect privacy) PIR protocols extend to this relaxed setting, up to constant factor losses. One of our bounds improves an earlier result of Goldreich et al. from the aforementioned paper. We also show that another of our bounds is essentially optimal by exhibiting an appropriate upper bound.

To explain our results in detail and compare them to previously known results, we begin with some necessarily definitions.

1.1 Preliminaries

We begin by introducing some notation. For an integer n , we let $[n]$ denote the set $\{1, 2, \dots, n\}$. For random variables X and Y that take values in the same set S , we write $X \approx_\delta Y$ to denote the fact that the L_1 distance (i.e., twice the variational distance) between the distributions of X and Y is at most δ . To be precise,

$$\sum_{a \in S} |\Pr[X = a] - \Pr[Y = a]| \leq \delta.$$

Definition 1.1. Let s, ℓ_q, ℓ_a be positive integers and ε, δ be reals in $[0, 1]$. An s -server $(\ell_q, \ell_a; \varepsilon, \delta)$ -PIR protocol $\mathcal{P}$ is a communication protocol between a single client, who holds an index $j \in [n]$, and s servers, each of whom holds a string $y \in \{0, 1\}^n$. Formally, $\mathcal{P}$ is specified by a triple (Q, A, Rec) of functions, where $Q : [s] \times [n] \times \{0, 1\}^\rho \rightarrow \{0, 1\}^{\ell_q}$, $A : [s] \times \{0, 1\}^n \times \{0, 1\}^{\ell_q} \rightarrow \{0, 1\}^{\ell_a}$, and $\text{Rec} : [n] \times \{0, 1\}^\rho \times (\{0, 1\}^{\ell_a})^s \rightarrow \{0, 1\}$ for some positive integer ρ . For compactness of notation, we shall write $Q_i(j, R)$ instead of $Q(i, j, R)$ and $A_i(y, z)$ instead of $A(i, y, z)$. Also, we shall drop the subscript on Q_i and A_i altogether when $s = 1$. The protocol operates as follows: the client generates a random string R distributed uniformly in $\{0, 1\}^\rho$ and, for each $i \in [s]$, sends a query string $Q_i(j, R)$ to server i . Upon receiving a query string z , server i sends an answer string $A_i(y, z)$ to the client. The client then outputs a recovered bit

$$\text{Out}(j, y, R) := \text{Rec}(j, R, A_1(y, Q_1(j, R)), \dots, A_s(y, Q_s(j, R))),$$

which is her guess at the value of y_j . The protocol must satisfy the following two conditions.

Correctness: $\forall j \in [n], y \in \{0, 1\}^n : \Pr_R[\text{Out}(j, y, R) = y_j] \geq 1 - \varepsilon$.

Privacy: $\forall i \in [s], j, k \in [n] : Q_i(j, R) \approx_\delta Q_i(k, R)$.

The parameter ℓ_q is called the query length, ℓ_a the answer length, ε the recovery error, and δ the privacy parameter of the protocol $\mathcal{P}$. The communication cost of $\mathcal{P}$ is $\text{cost}(\mathcal{P}) = s(\ell_q + \ell_a)$, the total number of bits communicated.

The goal in designing PIR protocols is to simultaneously reduce ε, δ , and $\text{cost}(\mathcal{P})$. We shall require that all servers receive queries of the same length and return answers of the same length. Since we only deal with constant values of s , this requirement causes no asymptotic loss.

When $\varepsilon = 0$, the protocol is said to have *perfect recovery*, and when $\delta = 0$, it is said to have *perfect privacy*. The bulk of theoretical research on PIR has focused on the case $\varepsilon = \delta = 0$. The work of Goldreich et al. [5] and that of Kerenidis and de Wolf [6] did consider the $\varepsilon > 0$ case. But relatively little attention has been paid to the $\delta > 0$ case, except for one result of Goldreich et al. mentioned below.

1.2 Our Results and Previous Work

We prove three lower bounds that allow $\delta > 0$. Let P be a 1-server $(\ell_q, \ell_a; \varepsilon, \delta)$ -PIR protocol. With the privacy requirement relaxed, even the 1-server case becomes nontrivial and it is not *a priori* clear that sublinear communication PIR is not possible. However, we show that for $\varepsilon = 0$, we must have $\text{cost}(P) \geq (1 - \delta/2)n = \Omega(n)$. We also show, via an upper bound, that this dependence to δ is essentially tight, up to terms quadratic in δ .

We also consider the more general case when both ε and δ can be nonzero. In this case, we show that $\text{cost}(P) \geq (1 - H(\varepsilon + \delta/2))n$ for sufficiently small ε and δ . Here H is the binary entropy function given by $H(x) := -x \lg x - (1-x) \lg(1-x)$; “lg” denotes logarithm to the base 2.

Finally, we consider 2-server schemes for nearly private information retrieval. It is known that, using two servers, $O(n^{1/3})$ communication can be achieved, even with $\varepsilon = \delta = 0$, via a number of different schemes; see, e.g., Chor et al. [4], Beimel, Ishai, and Malkin [3], and Woodruff and Yekhanin [11]. No strong general lower bound is known that comes close to matching this upper bound. However, a recent result of Razborov and Yekhanin [10] provides an $\Omega(n^{1/3})$ bound for protocols whose computation is restricted in a certain way. With arbitrary computations allowed, there *are* strong lower bounds known provided the answer length ℓ_a is short. The cleanest of these results are for the $\ell_a = 1$ case. In this case, Kerenidis and de Wolf [6] prove a lower bound of $(1 - H(11/14 - 4\varepsilon/7))n - 2$ on the communication cost when $\delta = 0$. Beigel, Fortnow, and Gasarch [2] prove a tight $n - 2$ lower bound when $\varepsilon = \delta = 0$.

Here, we prove a lower bound of $(1 - H(3/4 + 2\delta/3 - \sqrt{2\delta} - \varepsilon))n - 2$ when $\ell_a = 1$, for sufficiently small positive ε and δ . A lower bound handling positive ε and δ was proven by Goldreich et al. [5]. Their bound, for $\ell_a = 1$, is $(1 - 2\varepsilon - \delta)n/24 - 4$. (Note that our use of ε and δ is different from theirs; we have translated their bound into our notation.) To see that our bound is an improvement, consider the limiting case $\varepsilon \rightarrow 0, \delta \rightarrow 0$: our lower bound then approaches $0.19n - 2$, whereas the bound of [5] approaches $0.04n - 4$.

It is worth noting that the issue of lower bounds for PIR schemes with 3 or more servers has recently been largely settled, in a most dramatic way, by Yekhanin [12]: surprisingly low *upper* bounds hold.

2 Simple Upper Bounds

Here, we show simple improvements to the known upper bounds on the communication cost of PIR schemes by allowing imperfect privacy. As we shall see later,

the 1-server upper bound we obtain below is essentially optimal in the perfect recovery case.

Theorem 2.1. *For any $\delta > 0$, there is a PIR protocol with perfect recovery, privacy parameter δ , and communication cost at most $\lceil \lg n \rceil + \lceil (1 - \delta/(2 + \delta))n \rceil = (1 - \delta/2 + O(\delta^2))n + O(\log n)$.*

Proof. Let $\delta' = \delta/(2 + \delta)$. For each integer $j \in [n]$, define the sets

$$\begin{aligned} S_j &:= \{k \in [n] : 0 \leq (k - j) \bmod n \leq (1 - \delta')n\}, \\ T_j &:= \{k \in [n] : 0 \leq (j - k) \bmod n \leq (1 - \delta')n\}. \end{aligned}$$

It is important to keep in mind that $[n]$ denotes the set $\{1, 2, \dots, n\}$ whereas $x \bmod n$ takes values in $\{0, 1, \dots, n - 1\}$.

Design the function Q so that, when R is a uniform random string, $Q(j, R)$ is uniformly distributed on S_j . For $k \in [n]$ and $y \in \{0, 1\}^n$, let $A(y, k)$ return the concatenation, in some canonical order, of all y_j such that $j \in T_k$. It is easy to see that $k \in S_j \Leftrightarrow j \in T_k$; therefore $A(y, Q(j, R))$ is guaranteed to contain the desired bit y_j and we can design Rec so as to recover y_j from $Q(j, R)$ and $A(y, Q(j, R))$. Clearly, the PIR protocol given by (Q, A, Rec) has perfect recovery and communication cost at most $\lceil \lg n \rceil + |T_k| \leq \lceil \lg n \rceil + \lceil (1 - \delta')n \rceil$.

For all $j \in [n]$, we have $|S_j| \geq (1 - \delta')n$ and for $i \neq j$, we have $|S_i \setminus S_j| + |S_j \setminus S_i| \leq 2 \cdot |[n] \setminus S_j| \leq 2\delta'n$. Therefore, we can bound the protocol's privacy parameter as follows:

$$Q(i, R) \approx_{\delta''} Q(j, R), \quad \text{where } \delta'' \leq \frac{2\delta'n}{(1 - \delta')n} = \delta.$$

Thus, the protocol has all the desired properties.

A 2-Server Upper Bound. In a similar manner to the 1-server case, it is possible to add a δ -dependent coefficient to the $O(n^{1/3})$ upper bound for 2-server PIR. The idea is to suitably modify the covering codes scheme of Chor et al. [4]. The details are straightforward and hence omitted from this version.

3 1-Server Lower Bounds

3.1 Perfect Privacy and Recovery

Chor et al. [4] prove that, in the 1-server case with perfect privacy, n bits must be exchanged. Their argument goes as follows. A communication C (the string of exchanged bits) is said to be *possible* for (y, j) if there is a positive probability for C to happen when the database is y , and the user tries to obtain the j th bit. C is said to be *possible* for j if it is possible for some pair (y, j) . Let us fix a j and assume that the number of possible communications for j is less than 2^n . Then there exist different databases y, y' and C such that C is possible for both (y, j) and (y', j) . But by the privacy requirement, for every $k \in [n]$, C must also

be possible for (y, k) and (y', k) , since the queries are distributed equally, and the responses are determined by the queries. Pick an index j such that $y_j \neq y'_j$. We know that C is possible for both (y, j) and (y', j) , but C determines the output of the protocol, thus the protocol must yield the same bit, and we get a contradiction.

This argument fails in the almost secure case, since there is no requirement that the same communication be possible for all indices if it is possible for one. However, we can still obtain strong lower bounds, as we now show.

3.2 Nearly Private Schemes

Theorem 3.1. *Let $\mathcal{P}$ be a 1-server $(\ell_q, \ell_a; 0, \delta)$ -PIR protocol, where $\delta > 0$. Then $\ell_a \geq (1 - \delta/2)n$. In particular, $\text{cost}(\mathcal{P}) \geq (1 - \delta/2)n$.*

Proof. For $j \in [n]$ and $z \in \{0, 1\}^{\ell_q}$, let $p_{jz} = \Pr_R[Q(j, R) = z]$. Let $J_z = \{j : p_{jz} > 0\}$. It is easy to verify that

$$|J_z|p_{1z} \geq \sum_{j=1}^n \min\{p_{1z}, p_{jz}\} = \sum_{j=1}^n \left(\frac{p_{1z} + p_{jz}}{2} - \frac{|p_{1z} - p_{jz}|}{2} \right).$$

This implies

$$\begin{aligned} \sum_{z \in \{0, 1\}^{\ell_q}} |J_z|p_{1z} &\geq \sum_{j=1}^n \sum_{z \in \{0, 1\}^{\ell_q}} \left(\frac{p_{1z} + p_{jz}}{2} - \frac{|p_{1z} - p_{jz}|}{2} \right) \\ &\geq \sum_{j=1}^n \left(\frac{1+1}{2} - \frac{\delta}{2} \right) = (1 - \delta/2)n, \end{aligned}$$

where the final inequality follows from the privacy guarantee of $\mathcal{P}$. Since we have $\sum_{z \in \{0, 1\}^{\ell_q}} p_{1z} = 1$, there must exist a $z \in \{0, 1\}^{\ell_q}$ such that $|J_z| \geq (1 - \delta/2)n$. Fix such a z .

Suppose $\ell_a < |J_z|$. Let $Y := \{y \in \{0, 1\}^n : y_j = 0 \text{ for } j \notin J_z\}$. Then $|Y| = 2^{|J_z|}$. Meanwhile, the string $A(y, z)$ has length ℓ_a , so it lies in a set of size $2^{\ell_a} < 2^{|J_z|}$. By the pigeonhole principle, there exist distinct strings $y, y' \in Y$ such that $A(y, z) = A(y', z)$. Let j be an index such that $y_j \neq y'_j$. Then $j \in J_z$. Therefore, $p_{jz} > 0$, i.e., there exists an R such that $Q(j, R) = z$. Since $\mathcal{P}$ has perfect recovery, for this R we must have

$$y_j = \text{Rec}(j, R, A(y, z)) = \text{Rec}(j, R, A(y', z)) = y'_j,$$

which is a contradiction. This proves that $\ell_a \geq |J_z| \geq (1 - \delta/2)n$.

3.3 Nearly Private Schemes with Imperfect Recovery

We now turn to the imperfect recovery case. We prove our lower bound for this case by a reduction from a communication problem with a well known lower bound. Later, we use a much more sophisticated version of the same idea for a 2-server lower bound.

The problem INDEX_n is a communication problem involving two players: Alice, who holds an n -bit string $x = x_1x_2 \dots x_n$ (with each $x_i \in \{0, 1\}$), and Bob, who holds an index $i \in [n]$. A one-way communication protocol for this problem operates as follows: Alice sends Bob a message based on x after which Bob outputs his guess at the bit x_i . Both players may use a public random string in making their decisions, i.e., the protocol is allowed to be public coin. Abayev [1] proved the following sharp lower bound on the communication cost of such a protocol.

Fact 3.2. *Any public coin one-way communication protocol for INDEX_n with error at most ε must communicate at least $(1 - H(\varepsilon))n$ bits.*

Theorem 3.3. *Let ε and δ be positive reals with $\varepsilon + \delta/2 < 1/2$. Then any 1-server $(\ell_q, \ell_a; \varepsilon, \delta)$ -PIR protocol has $\ell_a \geq (1 - H(\varepsilon + \delta/2))n$. In particular, the communication cost of such a protocol is at least $(1 - H(\varepsilon + \delta/2))n$.*

Proof. Suppose P is a 1-server $(\ell_q, \ell_a; \varepsilon, \delta)$ -PIR protocol that uses ρ bits of randomness. Let $\mathcal{D}_{jz}$ denote the conditional distribution of R given that $Q(j, R) = z$ and let $\text{Gen} : [n] \times \{0, 1\}^{\ell_q} \times \{0, 1\}^{\rho'} \rightarrow \{0, 1\}^\rho$ be such that $\text{Gen}(j, z, R')$ is distributed according to $\mathcal{D}_{jz}$ when R' is distributed uniformly in $\{0, 1\}^{\rho'}$. Further, define $f : [n] \times \{0, 1\}^n \times \{0, 1\}^{\ell_q} \times \{0, 1\}^{\rho'} \rightarrow \{0, 1\}$ as follows.

$$f(j, y, z, r') := \begin{cases} 0, & \text{if } \text{Rec}(j, \text{Gen}(j, z, r'), A(y, z)) = y_j, \\ 1, & \text{otherwise.} \end{cases}$$

The correctness condition for P implies

$$\begin{aligned} \mathbb{E}_{R, R'}[f(j, y, Q(j, R), R')] &= \Pr_{R, R'}[\text{Rec}(j, \text{Gen}(j, Q(j, R), R'), A(y, Q(j, R))) \neq y_j] \\ &= \Pr_R[\text{Rec}(j, R, A(y, Q(j, R))) \neq y_j] \\ &\leq \varepsilon. \end{aligned}$$

Now, using the privacy condition $Q(j, R) \approx_\delta Q(1, R)$ and the fact that R and R' are independent, we have

$$\mathbb{E}_{R, R'}[f(j, y, Q(1, R), R')] \leq \varepsilon + \frac{\delta}{2}.$$

In other words, the following is a public coin one-way communication protocol for the problem INDEX_n , with error at most $\varepsilon + \delta/2$. Alice and Bob share a pair of random strings (R, R') distributed uniformly in $\{0, 1\}^\rho \times \{0, 1\}^{\rho'}$. Alice, upon receiving y , sends Bob the message $\mu := A(y, Q(1, R))$. Bob, upon receiving j and μ , outputs $\text{Rec}(j, \text{Gen}(j, Q(1, R), R'), \mu)$ as his guess at y_j . Clearly, this protocol has cost at most ℓ_a . By Fact 3.2, we have $\ell_a \geq (1 - H(\varepsilon + \delta/2))n$, which completes the proof.

4 2-Server Lower Bounds

We now turn to the case of 2-server PIR protocols. As mentioned earlier, much less is known about lower bounds for such protocols. In particular, the only strong lower bounds known for protocols that may make arbitrary computations are when the answer size is restricted to be quite small. In particular, there are

strong results known for the case of one-bit answers. Here, we prove an asymptotically optimal lower bound for the case of one-bit answers, with imperfect privacy allowed.

Our proof uses a *quantum computation* framework first used by Kerenidis and de Wolf [6]. Below, we quickly review the basics of quantum computation and communication and the Kerenidis - de Wolf framework and argument. We then show how to extend the framework to allow imperfect privacy. For an in-depth explanation of quantum computation we refer the reader to the textbooks by Nielsen and Chuang [9] and by Kitaev, Shen and Vyalı [7].

4.1 Quantum Communication

For our purposes, a quantum state is to be thought of as analogous to the classical notion of a probability distribution over fixed-length bit strings. A distribution over n -bit strings can be thought of a vector in $[0, 1]^{2^n}$ with unit ℓ_1 -norm. Analogously, an n -qubit state is a vector in $\mathbb{C}^{2^n} = (\mathbb{C}^2)^{\otimes n}$ with unit ℓ_2 -norm. We fix an orthonormal basis for the Hilbert space $(\mathbb{C}^2)^{\otimes n}$ and label the 2^n basis vectors (called basis states) by the 2^n n -bit strings: it is customary to use Dirac notation and denote the vector labeled by the string a as $|a\rangle$. It is also customary to write, e.g., $|5\rangle$ for the 3-qubit state $|101\rangle$ because “101” is the binary representation of 5.

An n -qubit quantum state can evolve by the application of a unitary transformation in $U(2^n)$. It can also be measured in a variety of ways whose details need not concern us here. For our purposes, we need only consider the following type of measurement. Suppose we have a decomposition $(\mathbb{C}^2)^{\otimes n} = \mathcal{W}_1 \oplus \mathcal{W}_2 \oplus \dots \oplus \mathcal{W}_k$, and suppose W_j denotes the projection onto $\mathcal{W}_j$. Then we can measure an n -qubit state $|\phi\rangle$ according to this decomposition: we will obtain a random outcome in the set $[k]$, with the probability of outcome j being $\|W_j|\phi\rangle\|_2^2 = \langle\phi|W_j|\phi\rangle$.

A quantum communication protocol is like a (classical) communication protocol except that the communicating parties may send qubits (i.e., quantum states) to each other. The communication cost of a protocol is the number of qubits sent.

4.2 Perfect Privacy

Kerenidis and de Wolf prove a number of communication lower bounds for 2-server PIR schemes. However, their arguments only handle the perfect privacy case, although they do handle imperfect recovery. Their arguments are cast in a quantum communication framework whose key observation can be expressed thus: “a single quantum query can simulate two classical queries.” Using this observation, they build a 1-server “quantum PIR scheme” and then prove lower bounds on its communication in a way analogous to our 1-server lower bounds. In particular, the appropriate quantum analog of Abılayev’s lower bound (Fact 3.2) turns out to be a lower bound for quantum random access codes, due to Nayak [8].

We now outline Kerenidis and de Wolf’s argument, using our own terminology. We find it convenient to remove the intermediate steps of a quantum PIR

scheme and a quantum random access code; instead, we show that a 2-server PIR scheme with good enough parameters implies a one-way *quantum* communication protocol for INDEX_n with low communication cost. The desired PIR lower bound then follows from the aforementioned result of Nayak [8], which can be restated thus.

Fact 4.1. *A one-way quantum communication protocol for INDEX_n with error probability ε must communicate at least $(1 - H(\varepsilon))n$ qubits.*

We now fill in some details. Suppose P is a 2-server $(\ell_q, 1; \varepsilon, \delta)$ -PIR protocol, given by (Q, A, Rec) , that uses ρ bits of randomness. We associate with P a certain collection $\{|\phi_{jy}\rangle\}$ of $(\rho + 4 + \ell_a)$ -qubit quantum states. To define these, we use the basis states $|r, i, i, z\rangle : r \in \{0, 1\}^\rho, i \in \{0, 1, 2\}, z \in \{0, 1\}^{\ell_q}$. We set $c := 1/\sqrt{3 \cdot 2^\rho}$ and, for notational convenience, we define $Q_0(j, r) = 0^{\ell_q}$ and $A_0(y, z) = 0$ for all $j \in [n], r \in \{0, 1\}^\rho, y \in \{0, 1\}^n$ and $z \in \{0, 1\}^{\ell_q}$. Also, for $(i, j, z) \in \{0, 1, 2\} \times [n] \times \{0, 1\}^{\ell_q}$, we define the set $S_{ijz} := \{r \in \{0, 1\}^\rho : Q_i(j, r) = z\}$. Finally, we define $|\phi_{jy}\rangle$ as follows:

$$|\phi_{jy}\rangle := \sum_{r \in \{0, 1\}^\rho} c |r\rangle \left(|0, 0, 0^{\ell_q}\rangle + (-1)^{A_1(y, Q_1(j, r))} |1, 1, Q_1(j, r)\rangle + (-1)^{A_2(y, Q_2(j, r))} |2, 2, Q_2(j, r)\rangle \right).$$

The significance of this quantum state is brought out by the following fact, implicit in the work of Kerenidis and de Wolf.

Fact 4.2 (Kerenidis and de Wolf [6]). *By measuring $|\phi_{jy}\rangle$ appropriately, one can obtain a random 2-bit outcome (β_1, β_2) such that*

$$\Pr[(\beta_1, \beta_2) = (A_1(y, Q_1(j, r)), A_2(y, Q_2(j, r)))] \geq 3/4.$$

Therefore, by applying the function Rec to the measured outcome, one can obtain a bit that equals y_j with probability at least $3/4 - \varepsilon$. In fact, the probability of correctly recovering y_j can be further improved to $11/14 - 4\varepsilon/7$ by using a (classical) postprocessing trick.

To see how this fact can be used to obtain the desired communication protocol, note that

$$\begin{aligned} |\phi_{jy}\rangle &= \sum_{r \in \{0, 1\}^\rho} \sum_{i=0}^2 (-1)^{A_i(y, Q_i(j, r))} c |r, i, i, Q_i(j, r)\rangle \\ &= \sum_{i=0}^2 \sum_{z \in \{0, 1\}^{\ell_q}} \sum_{r \in S_{ijz}} (-1)^{A_i(y, z)} c |r, i, i, z\rangle \\ &= \sum_{i=0}^2 \sum_{z \in \{0, 1\}^{\ell_q}} |\chi_{ijz}\rangle \cdot (-1)^{A_i(y, z)} c \sqrt{|S_{ijz}|} |i, z\rangle, \end{aligned}$$

where $|\chi_{ijz}\rangle := |S_{ijz}|^{-1/2} \sum_{r \in S_{ijz}} |r, i\rangle$. Let U_j be a unitary transformation that maps $|0^\rho, 0, i, z\rangle$ to $|\chi_{ijz}\rangle|i, z\rangle$. The protocol for INDEX_n works as follows. Alice, on input y , prepares the quantum state

$$|\psi_{jy}\rangle := \sum_{i=0}^2 \sum_{z \in \{0,1\}^{\ell_q}} (-1)^{A_i(y,z)} c \sqrt{|S_{ijz}|} |i, z\rangle \quad (1)$$

and sends it to Bob. Although it seems at first glance that $|\psi_{jy}\rangle$ depends on j , it in fact doesn't, because the perfect privacy guarantee of P implies that for $j, k \in [n]$,

$$\frac{|S_{ijz}|}{2^\rho} = \Pr_R[Q_i(j, R) = z] = \Pr_R[Q_i(k, R) = z] = \frac{|S_{ikz}|}{2^\rho}. \quad (2)$$

Bob, upon receiving $|\psi_{jy}\rangle$, constructs the state $|0^\rho, 0\rangle|\psi_{jy}\rangle$ using ρ qubits of his own and applies U_j to it. By definition of U_j , the state that Bob obtains is $|\phi_{jy}\rangle$. He then uses the procedure implied by Fact 4.2 to compute his output bit, which is correct with probability at least $11/14 - 4\varepsilon/7$. Since $|\psi_{jy}\rangle$ is a $(2 + \ell_q)$ -qubit state, the communication cost of this protocol is $2 + \ell_q$. Fact 4.1 now implies that $\ell_q \geq (1 - H(11/14 - 4\varepsilon/7))n - 2$, giving us a lower bound on $\text{cost}(P)$.

4.3 The Nearly Private Case

Without perfect privacy, the argument above does not work. This is because Eq. (2) no longer holds, which makes the above quantum communication protocol ill-defined: Alice can no longer prepare the state $|\psi_{jy}\rangle$ because it *might* depend on j , which Alice does not know. However, we shall show that Alice can get away with sending Bob the state $|\psi_{1y}\rangle$, provided a sufficiently strong privacy guarantee holds.

Theorem 4.3. *Let ε and δ be sufficiently small positive reals. Then any 2-server $(\ell_q, 1; \varepsilon, \delta)$ -PIR protocol has $\ell_q \geq (1 - H(3/4 + 2\delta/3 - \sqrt{2\delta - \varepsilon}))n - 2$. In particular, the communication cost of such a protocol is at least $\Omega_{\varepsilon, \delta}(n)$.*

Proof. We use the framework and notation of Section 4.2. Suppose P is a 2-server $(\ell_q, 1; \varepsilon, \delta)$ -PIR protocol. Consider the following one-way communication protocol for INDEX_n : Alice, on input y , sends Bob the $(2 + \ell_q)$ -qubit quantum state $|\psi_{1y}\rangle$. Bob, upon receiving it, constructs the state $|0^\rho\rangle|\psi_{1y}\rangle$ defined by Eq. (1) and applies the unitary transformation U_j to it. He then measures the resulting state $|\phi'_{jy}\rangle$ as mentioned in Fact 4.2.

Let us eschew the additional “11/14 trick” referred to in Fact 4.2 and instead consider the probability p that Bob obtains the “correct” outcome — i.e., the pair of bits $(A_1(y, Q_1(j, r)), A_2(y, Q_2(j, r)))$ — when he uses the same measurement on the state $|\phi'_{jy}\rangle$. Let W be the projection operator corresponding to the desired outcome, so that $\|W|\phi_{jy}\rangle\|_2^2 \geq 3/4$ and $p = \|W|\phi'_{jy}\rangle\|_2^2$. Then

$$\|W|\phi'_{jy}\rangle\|_2 \geq \|W|\phi_{jy}\rangle\|_2 - \|W(|\phi_{jy}\rangle - |\phi'_{jy}\rangle)\|_2 \geq \frac{\sqrt{3}}{2} - \| |\phi_{jy}\rangle - |\phi'_{jy}\rangle \|_2. \quad (3)$$

Now,

$$\begin{aligned} \|\phi_{jy}\rangle - \phi'_{jy}\rangle\|_2^2 &= \|\lvert 0^\rho, 0\rangle\phi_{jy}\rangle - \lvert 0^\rho, 0\rangle\phi'_{jy}\rangle\|_2^2 \\ &= \|\lvert \psi_{jy}\rangle - \lvert \psi_{1y}\rangle\|_2^2 \end{aligned} \quad (4)$$

$$\begin{aligned} &= \sum_{i=0}^2 \sum_{z \in \{0,1\}^{\ell_q}} c^2 \left(\sqrt{|S_{ijz}|} - \sqrt{|S_{i1z}|} \right)^2 \\ &\leq \sum_{i=0}^2 \sum_{z \in \{0,1\}^{\ell_q}} c^2 \left| |S_{ijz}| - |S_{i1z}| \right| \end{aligned} \quad (5)$$

where Eq. (4) holds because U_j is unitary, and Eq. (5) is obtained by invoking Eq. (1). Since P has privacy parameter δ , for $i \in \{1, 2\}$ we have $\sum_{z \in \{0,1\}^\rho} \left| |S_{ijz}| - |S_{i1z}| \right| \leq 2^\rho \delta$. Also, by design, $S_{0jz} = S_{01z}$ for all z . Putting these facts together and using Eq. (3) gives

$$p = \|W|\phi'_{jy}\rangle\|_2^2 \geq \left(\frac{\sqrt{3}}{2} - \sqrt{2c^2 2^\rho \delta} \right)^2 = \frac{3}{4} + \frac{2\delta}{3} - \sqrt{2\delta}.$$

Since Bob eschews the classical postprocessing (the “11/14 trick”), the probability that he correctly outputs y_j is at least the above quantity minus the probability that the PIR scheme errs, i.e., at least $3/4 + 2\delta/3 - \sqrt{2\delta} - \varepsilon$. The theorem follows.

5 Conclusion

We have found that, in the 1-server case and in the binary 2-server case, relaxing the privacy requirements on a private information retrieval (PIR) scheme by allowing it to leak a small amount of information about the client’s index does not allow more than a constant factor improvement in the communication cost. The question of whether improvements can be obtained for the general 2-server case remains open.

References

1. Abloyev, F.: Lower bounds for one-way probabilistic communication complexity and their application to space complexity. *Theoretical Computer Science* 175(2), 139–159 (1996)
2. Beigel, R., Fortnow, L., Gasarch, W.: A tight lower bound for restricted PIR protocols. *Comput. Complexity* 15(1), 82–91 (2006)
3. Beimel, A., Ishai, Y., Malkin, T.: Reducing the servers computation in private information retrieval: PIR with preprocessing. In: Bellare, M. (ed.) *CRYPTO 2000*. LNCS, vol. 1880, pp. 56–74. Springer, Heidelberg (2000)
4. Chor, B., Goldreich, O., Kushilevitz, E., Sudan, M.: Private information retrieval. *J. ACM* 45(6), 965–982 (1998)

5. Goldreich, O., Karloff, H., Schulman, L., Trevisan, L.: Lower bounds for linear locally decodable codes and private information retrieval. In: Proc. 17th Annual IEEE Conference on Computational Complexity, pp. 175–183. IEEE Computer Society Press, Los Alamitos (2002)
6. Kerenidis, I., de Wolf, R.: Exponential lower bound for 2-query locally decodable codes. *J. Comput. Syst. Sci.* 69(3), 395–420 (2004) (Preliminary version in Proc. 35th Annual ACM Symposium on the Theory of Computing, pp.106–115 (2003))
7. Kitaev, A.Y., Shen, A.H., Vyalii, M.N.: *Classical and Quantum Computation*. American Mathematical Society (2002)
8. Nayak, A.: Optimal lower bounds for quantum automata and random access codes. In: Proc. 40th Annual IEEE Symposium on Foundations of Computer Science, pp. 124–133. IEEE Computer Society Press, Los Alamitos (1999)
9. Nielsen, M.A., Chuang, I.L.: *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge (2000)
10. Razborov, A., Yekhanin, S.: An $\Omega(n^{1/3})$ lower bound for bilinear group based private information retrieval. In: Proc. 47th Annual IEEE Symposium on Foundations of Computer Science, pp. 739–748. IEEE Computer Society Press, Los Alamitos (2006)
11. Woodruff, D., Yekhanin, S.: Towards 3-query locally decodable codes of subexponential length. In: Proc. 20th Annual IEEE Conference on Computational Complexity, 2005, pp. 275–284. IEEE Computer Society Press, Los Alamitos (2005)
12. Yekhanin, S.: Towards 3-query locally decodable codes of subexponential length. In: Proc. 39th Annual ACM Symposium on the Theory of Computing, 2007, ACM Press, New York (to appear, 2007)